


LQCD Project Review Response



Germantown, MD

Aug 8, 2005



- Don Holmgren, FNAL, Project Mgr.
- Bakul Banerjee, FNAL, Assoc. Prj. Mgr.
- Chip Watson, JLab, Site Mgr.
- Tom Schlagel, BNL, Site Mgr.
- Bob Sugar, UCSB, Chairman LQCD Executive Committee



Agenda

- Exhibit 300 discussions
 - Project scope
 - Manpower: project and base
 - Power/cooling, storage
- FY06 plans
 - Initial conditions
 - Scientific requirements
 - Proposal
- Review response (as time permits)




Exhibit 300 Discussions



Project Scope

- The construction/purchase of new clusters/systems
- The operation of the U.S. QCDOC, the existing SciDAC clusters, and the systems acquired during the project
- Labor for system administration, hardware support, and site management will be paid by the project
- All labor for scientific software support as well as support of the scientific needs of users will be paid by laboratory base contributions and by SciDAC project
- Software development is out of scope




Staffing – FY06

Project + Base/SciDAC FTE	sysadmin / technician	Scientific software & user support	site management
BNL	0.75	0.5	0.25
FNAL	1.75	0.5	0.25
JLab	0.65	0.5	0.25

Manpower increases at FNAL, JLab after FY06 as size of operated systems grows.


Staffing Costs

LQCD Computing FTE Costs



Power/Cooling

- See spreadsheet for details
- Linear power growth fits 2001-2005 LQCD clusters
 - New Intel processor in 2006: 30 Watts compared with current model's ~100 Watts (60 Watts system power)





Power/Cooling

- Costs shown use linear power growth model
 - Alternate site scenario (FNAL/JLab, FNAL, JLab, FNAL)
 - Exponential model doubles power costs

Cost (\$k)	FY06	FY07	FY08	FY09
BNL	\$80.3	\$83.5	\$86.8	\$90.3
FNAL	\$150.0	\$319.9	\$420.0	\$400.0
JLAB	\$106.1	\$110.4	\$182.7	\$288.8



Tape Storage

- Gauge Configurations
 - Permanent, redundant storage
 - FY06: 9.0 TB total, 45 tapes, \$3.3K tape value for a single copy
 - Laboratories are contributing tapes
 - Out years: tape densities increase
 - But, gauge configuration volume will grow in size
 - Have budgeted appropriately




Tape Storage

- Propagators and Matrix Elements
 - Temporary storage: 6-18 months depending upon calculation campaign strategies
 - FY06: ~ 142 TB, ~ 710 tapes
 - FNAL is migrating existing experimental data tapes to higher density drives (50 GB to 200 GB)
 - At FNAL this frees more than sufficient # tapes for LQCD
 - At JLab and BNL, sufficient tapes (~ 200) remain from FY05 for FY06
 - Labs will contribute tapes for FY06:
 - Tape value: \$52.7K
 - Tape silo "slot" value: \$39.2K (\$55 per slot)

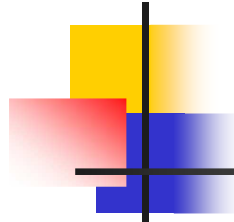


Exhibit 300 Questions

- Should base contributions be shown? If so, how?
- When are revisions due?
- Changes that need to be done:
 - Modifications to security section
 - Revised budget breakdown and corresponding deliverables
 - Minor PBMS modifications
- Program Manager certification level



FY06 Plans



FY06 Plans – Initial Conditions

- Deployed hardware as of Oct 1, 2005:
 - BNL: 4.2 TF U.S. QCDOC
 - FNAL: 0.14 TF Myrinet cluster
0.86 TF Infiniband cluster
 - JLab: 0.19 TF GigE mesh cluster
0.46 TF GigE mesh cluster



Infiniband Results to Date

- SciDAC Infiniband clusters (all FNAL):
 - 32-node prototype (PCI), 2004
 - 520-node production (PCI Express), 2005
- Very successful
 - < \$1/Mflop for Domain Wall
 - Scalable to thousands of nodes
- But:
 - Difficult integration (easier than Myrinet at same age)
 - Expertise is won only through experience, working with vendor – poor written documentation
 - Experience within project is only at FNAL




FY06 Proposal

- It is clear that Intel/AMD + Infiniband are the most cost effective choice for FY06/FY07
- To minimize risk, the project must develop Infiniband cluster expertise at JLab
 - Final year SciDAC funds can be used to start the first JLab Infiniband cluster
 - Project funds should be used to extend this cluster to the size that meets scientific & project requirements
- A capability cluster of order 2 Tflops will be deployed at FNAL



Scientific Requirements

- JLab analysis cluster: 0.4 Tflops (DWF)
 - Sufficient capability for DWF algorithm development (an approved project)
 - Sufficient capability for analysis of DWF quarks using large asqtad lattices (also an approved project)
 - JLab cluster should be 256 nodes
 - Based on FNAL FY05 cluster: 1.65 GF/node à 0.42 TF
 - At 1/4 scale of assumed FY07+ cluster sizes, this FY06 cluster will provide the necessary experience at sufficient scale
 - 128 nodes purchased via FY05 SciDAC / FY06 base in Fall, 128 identical nodes purchased via Project in FY06
 - Project nodes purchased via exercise of option on FY05 SciDAC procurement – no separate procurement necessary




Scientific Requirements

- FNAL Capability Cluster: 2+ Tflops
 - $40^3 \times 96$ and $48^3 \times 128$ asqtad lattices require minimum 1 Tflop partitions
 - With FY05 SciDAC contributions (\$150K), estimated cluster size is 940 nodes, 2.45 GF/node
 - Cluster will be used for 1 or 2 configuration generation streams plus analysis streams
 - Single FY06 procurement



Timing

- To reach 2.45 Gflop/node performance, will need to use Intel 1066 MHz FSB, faster processors, FBD memory technology
 - Intel: calendar 1st quarter à in channel for bidding by February?
 - February 1st RFP à production May 1
 - Dependency on FNAL construction
 - Swing space is available if beneficial occupancy delayed; construction approved and procurement activities have started
 - Beneficial occupancy: April 2



Review Report Response



Significance and Merit of the Proposed Initiative

- Recommendation 1: “[F]acilitate exploratory studies in algorithms and comparative quantum field theory by allocating some time on the facility to this type of project.”
- Response: This has been a long term scientific goal, and we will continue to allocate time for such studies. We have a very promising collaborative effort with the TOPS ISIC (David Keyes, adaptive multigrid) as part of our SciDAC work, and we will propose specific support for algorithm development in our upcoming SciDAC II proposal.



Significance and Merit of the Proposed Initiative

- Recommendation 2: “The team should develop a plan to incorporate specific visualization goals and approaches, as well as ensure sufficient visualization resources to make the approach feasible.”
- Response: Software development is not within the scope of this project. However, the SciDAC project plans to address this area. Further, as appropriate, acquisition plans will address visualization needs, providing the necessary hardware and (likely commercial) software infrastructure.



Significance and Merit of the Proposed Initiative

- Recommendation 3: “It is vital to the long-term health of the subject that young researchers get attracted into it. The team should consider ways in which this facility can be used to help the development of young researchers.”




Significance and Merit of the Proposed Initiative

- **Response:** We will continue to give high priority to proposals for computer time by young researchers. We will push to create new faculty positions and laboratory staff positions in our field, including joint appointments between the host laboratories and universities. A recent example was the appointment of Kostas Orginos, an outstanding young lattice gauge theorist, to a tenure track position by William and Mary/JLab. We plan to organize a series of summer schools in lattice gauge theory for graduate postdoctoral students. The Institute for Nuclear Physics in Seattle has agreed to host a summer school in 2007.



Status of the Technical Design

- Finding 1: “The projected budgets and anticipated Moore’s Law improvements in computational power should allow for the yearly acquisition of new clusters at about the same delivered performance on LQCD applications as the aggregate of existing computing resources.”
- Comment: In FY2006, the project begins with 5.8 Tflops of existing capacity and will add approximately 2.75 Tflops of new capacity. There are insufficient funds in any year of the project to add hardware matching the existing aggregate capacity; rather, roughly 25-30% additional capacity will be added each year.



Status of the Technical Design

- Recommendation 1: "...acquisition plan [should] be modified to allow for a single joint acquisition, possibly every other year, alternating between JLab and FNAL that would allow the delivery of resources to the program promptly in FY06 and beyond. The number of procurements should be reduced from eight to three or four."



Status of the Technical Design

- **Response:** We agree that 8 procurements should be reduced to 3 or 4. Procurements will be a collaborative effort of the Project Manager and the Site Managers. In FY06, we propose that the cluster designed by the project be procured by FNAL. The project strongly feels that the cluster should be housed at FNAL because of their experience with Infiniband fabrics. The project also feels that it is critical that JLab gain experience with Infiniband, and recommends that JLab procure a 128-node cluster in FY05 with SciDAC and FY06 base funds, and extend this cluster to 256-nodes in FY06 with project funds; the resulting 400 Gflop cluster will meet the scientific needs of the approved DWF algorithm development and analysis of DWF quarks on asqtad lattices.




Status of the Technical Design

- Response (con't): In subsequent years the project will select the hardware (clusters vs. other supercomputers) and the location of the hardware in order to maximize the science according to the planned scientific program for the following year(s).




Status of the Technical Design

- Recommendation 2: “If the revised FNAL schedule is accurate ... the team should decide on the site for the [FY2006] computer based on where it can deliver the most science for the dollars invested.
- Response: FNAL has committed to a schedule for the computer room refurbishment which will allow beneficial occupancy by April 2, 2006. We will follow the Program Manager’s advice regarding the timing of the Federal Budget and will schedule the release of the RFP to first commit funds in fiscal Q2. Hardware delivery would therefore match the FNAL construction schedule. Further, we note that Intel roadmaps strongly favor delaying the procurement until mid-Q2.



Status of the Technical Design

- Recommendation 3: “The cluster integration plan should be written down and an architectural diagram ... clearly indicated. The plan should also include the software development and integration work items necessary to bring these resources into production. This plan should be presented to the LQCD scientific advisory board for review and approval.”



Status of the Technical Design

- **Response:** During the SciDAC project cluster designs were reviewed by the Oversight Committee, which included computing experts from outside LQCD. We will continue to follow this procedure and will also obtain the input of the LQCD Executive Committee for each plan; this committee will have the responsibility of certifying that the plans fully meet the scientific requirements. The project plans will include the requested architectural diagrams as well as software development and integration details.



Status of the Technical Design

- Recommendation 4: “The LQCD project plan should be expanded to identify dependencies on SciDAC and other projects for technology necessary for building the Metafacility.”
- Response: The WBS will be expanded appropriately, with Level 1 and/or Level 2 deliverables and milestones. We note that the risks are moderate (SciDAC) to high (ILDG and other GRID developments), since the necessary software comes from projects external to this one; however, the consequences to the deliverables of schedule slip for GRID software are minimal and manageable.



Feasibility/Completeness of Proposed Budget and Schedule

- Recommendation 1: “The project should consider alternative deployment strategies that result in fewer, larger systems over the same time period”.
- Response: In FY06, a SciDAC Infiniband cluster at JLab similar to the FNAL FY05 cluster will be expanded, and a large FNAL cluster will be procured. In the subsequent 3 years, 2 to 3 additional large procurements will occur, depending upon the timing of introductions of improved hardware to the market.



Feasibility/Completeness of Proposed Budget and Schedule

- Recommendation 2: “The project should provide a cost benefit analysis for one site, two sites and three sites as part of the planning.”
- Response: The project will perform and include this analysis in the project plans.



Feasibility/Completeness of Proposed Budget and Schedule

- Recommendation 3: “The cost projections for storage and consumables should be done to the same level as the costs for computational resources in order to ensure the user requirements are met in a balanced manner.”
- Response: We have gathered much additional information about the quantity and lifetime of data products and have modified the cost projections accordingly. The propagators discussed at the review, which take up most of the required storage, are intermediate data products that can be deleted 12-18 months after generation.



Feasibility/Completeness of Proposed Budget and Schedule

- Recommendation 4: “The team should ensure wide impact of the valuable SciDAC-funded prototyping work with more timely publication of their results ... on the web site ... [and] in more widely shared publications and conferences. This effort should also seek out collaborations with other architectural and performance evaluation efforts”.



Feasibility/Completeness of Proposed Budget and Schedule

- **Response:** The project will increase the number of presentations and publications as recommended. We will also widen our collaborative efforts.



Feasibility/Completeness of Proposed Budget and Schedule

- Recommendation 5: “The project team should reevaluate the principles used to determine which costs are included within the project to ensure an accurate presentation of the overall cost of the effort to the DOE.”
- Response: The project has now produced effort and cost breakdowns showing the contributions from the laboratory base budgets, including power costs. We have also corrected the inconsistencies present in the earlier project plan of the lab/base effort breakdowns between the three labs.



Relevance of Prototyping Efforts and the Status and Plans for Developing the Required Software.

- Recommendation 1: “The LQCD team should continue to monitor the market and benchmark the available options. The team should ... develop an integrated prototyping activity. [T]he committee recommends the use of vendor-provided drivers to increase the communications performance (e.g., QMP over VAPI instead of MPI).”



Relevance of Prototyping Efforts and the Status and Plans for Developing the Required Software.

- **Response:** In FY05/06, year 5 SciDAC funds will continue to be used for benchmarking available options. We will integrate this effort between FNAL and JLab, and as appropriate (e.g., storage hardware) with BNL. After FY06, SciDAC II funds, or base funds, will contribute to these activities, as they are of great interest and utility to the host laboratories. We will devote SciDAC-supported effort to implementation of the communications library (QMP) over VAPI; this work has begun.



Relevance of Prototyping Efforts and the Status and Plans for Developing the Required Software.

- Recommendation 2: “SciDAC support has been, and continues to be, absolutely essential for the success of the LQCD project. [T]he committee recommends that the DOE consider continuing these efforts.”
- Response: The project enthusiastically concurs. Those project members involved in preparation of SciDAC II proposals will cite this recommendation.



Effectiveness of the Proposed Management Structure

- Recommendation 1: “Operations agreements with the sites over the lifetime of the project should be executed, which cover all contributions that are not included in the project cost...”
- Response: We are preparing memoranda of understanding with the host laboratories. All such agreements will be referenced in our Project Execution Plan.



Effectiveness of the Proposed Management Structure

- Recommendation 2: “The project should develop, and update on a yearly basis, a project-wide system deployment plan that optimizes the opportunity to deliver new science without artificial constraints on which programs can fund work at the three partner laboratories.”
- Response: The project has developed plans which are not constrained as noted in the recommendation.



Effectiveness of the Proposed Management Structure

- Recommendation 3: “The strongly site-based management scheme reflects in part the history of forming this project. The laboratories should integrate their planning, prototyping and procurement activities. The approach to this should be in the revised PEP.
- Response: Prototyping and procurement activities will be integrated and the approach described in the revised PEP.



Effectiveness of the Proposed Management Structure

- Recommendation 4: “Ensure the WBS is a tool for integrated planning as well as integrated reporting. The reporting should also document the actual physics output measured in terms of the allocations made by the Scientific Program Committee.”



Effectiveness of the Proposed Management Structure

- Recommendation 4.1: “Incorporate schedules for integrated review of outyear plans into WBS to occur no later than June preceding beginning of FY.”
- Response: This will be done for FY07, FY08, and FY09.



Effectiveness of the Proposed Management Structure

- Recommendation 4.2: “Expand procurement processes to include all three sites and possibly external experts, including evaluation of joint procurements.”
- Response: Procurements processes will include all three sites and when possible will be executed by a single site. Evaluation of proposed procurements will include outside experts (Oversight Committee).



Effectiveness of the Proposed Management Structure

- Recommendation 4.3: “Consider integrating technology tracking, hardware and software prototyping across all three sites.”
- Response: Cluster hardware and software prototyping will be integrated across JLab and FNAL. To the extent feasible, storage related hardware and software prototyping will be integrated across all three sites; however, we note that at BNL the QCDOC currently only supports NFS-based disk systems. FNAL and JLab plan to use the same parallel filesystem (dCache or an alternative) and will integrate the related hardware and software efforts.



Effectiveness of the Proposed Management Structure

- Recommendation 4.4: “Since there are strong dependencies on some external efforts (SciDAC, ILDG, FNAL construction) schedule and contingency for these need to be in the WBS.”
- Response: We will include all such dependencies in the WBS and in the risk management strategy.



Effectiveness of the Proposed Management Structure

- Recommendation 4.5: “The laboratories should report the monthly progress of each laboratory in providing the capabilities and capacity agreed to by the Scientific Program Committee.”
- Response: The monthly reporting by the project will include this information as well as the actual monthly physics output (delivered flops by scientific project) at each site.



Effectiveness of the Proposed Management Structure

- Recommendation 5: “Consider moving metafacility operations to the integrated project office.”
- Response: This will be done. In the organizational chart, the Metafacility Operations Manager (MFO) will report to the Project Manager; however, accounting will be done at JLab, where the MFO resides, similar to the site managers.



Effectiveness of the Proposed Management Structure

- Recommendation 6: “Charters for Executive and Scientific Program Committees including how members are chosen should be produced and included in the PEP.
- Response: This will be done.



Effectiveness of the Proposed Management Structure

- Recommendation 7: “Try to integrate the Change Control Board and Scientific Program Committee review of change proposals.”
- Response: The Executive Committee will establish a permanent CCB for the duration of the project with at least three members. Members will be appointed by the Chair of the LQCD Executive Committee from the Executive Committee membership and from the user community. These members will not be a part of the participating Laboratories.